

PATENT APPLICATION

Systematic Encoding and Decoding of Chain Reaction Codes

Inventors

M. Amin Shokrollahi, a citizen of Germany,
residing at 5780 Chandler Court
San Jose, California 95123
United States of America

Michael G. Luby, a citizen of the United States,
residing at 1133 Miller Avenue
Berkeley, California 94708
United States of America

Assignee

Digital Fountain, Inc.
39141 Civic Center Drive
Suite 300
Fremont, California 94538
United States of America

Representative

Clifford B. Perry
732 North Coast Hwy 101
Suite B
Encinitas, California 92024
United States of America

Cross-Reference to Related Applications

[0001] This application claims the benefit of US Provisional Application no. 60/319,597 entitled "Systematic Encoding and Decoding of Chain Reaction Codes," filed October 5, 2002, the contents of which are herein incorporated by reference in its entirety for all purposes.

Background

[0002] The following invention relates to systems and methods for encoding and decoding data of all types, and more particularly to systems and methods for encoding and decoding data using chain reaction codes.

[0003] Transmission of data between a sender and a recipient over a communications channel has been the subject of much literature. Preferably, but not exclusively, a recipient desires to receive an exact copy of data transmitted over a channel by a sender with some level of certainty. Where the channel does not have perfect fidelity (which covers most of all physically realizable systems), one concern is how to deal with data lost or garbled in transmission. Lost data (erasures) are often easier to deal with than corrupted data (errors) because the recipient cannot always tell when corrupted data is data received in error. Many error-correcting codes have been developed to correct for erasures and/or for errors. Typically, the particular code used is chosen based on some information about the infidelities of the channel through which the data is being transmitted and the nature of the data being transmitted. For example, where the channel is known to have long periods of infidelity, a burst error code might be best suited for that application. Where only short, infrequent errors are expected a simple parity code might be best.

[0004] Another consideration in selecting a code is the protocol used for transmission. In the case of the Internet, a packet protocol is used for data transport. That protocol is called the Internet Protocol or "IP" for short. When a file or other block of data is to be transmitted over an IP network, it is partitioned into equal size input symbols and input

symbols are placed into consecutive packets. The “size” of an input symbol can be measured in bits, whether or not the input symbol is actually broken into a bit stream, where an input symbol has a size of M bits when the input symbol is selected from an alphabet of 2^M symbols. In such a packet-based communication system, a packet oriented coding scheme might be suitable.

[0005] A transmission is called reliable if it allows the intended recipient to recover an exact copy of the original file even in the face of erasures in the network. On the Internet, packet loss often occurs because sporadic congestion causes the buffering mechanism in a router to reach its capacity, forcing it to drop incoming packets. Protection against erasures during transport has been the subject of much study.

[0006] The Transport Control Protocol (“TCP”) is a point-to-point packet control scheme in common use that has an acknowledgment mechanism. Using TCP, a sender transmits ordered packets and the recipient acknowledges receipt of each packet. If a packet is lost, no acknowledgment will be sent to the sender and the sender will resend the packet. With protocols such as TCP, the acknowledgment paradigm allows packets to be lost without total failure, since lost packets can just be retransmitted, either in response to a lack of acknowledgment or in response to an explicit request from the recipient.

[0007] Although acknowledgment-based protocols are generally suitable for many applications and are in fact widely used over the current Internet, they are inefficient, and sometimes completely infeasible, for certain applications, as described in Luby I.

[0008] One solution that has been proposed to solve the transmission problem is to avoid the use of an acknowledgment-based protocol, and instead use Forward Error-Correction (FEC) codes, such as Reed-Solomon codes, Tornado codes, or chain reaction codes, to increase reliability. The basic idea is to send output symbols generated from the content instead of just the input symbols that constitute the content. Traditional erasure correcting codes, such as Reed-Solomon or Tornado codes, generate a fixed number of output symbols for a fixed length content. For example, for K input symbols, N output

symbols might be generated. These N output symbols may comprise the K original input symbols and $N-K$ redundant symbols. If storage permits, then the server can compute the set of output symbols for each content only once and transmit the output symbols using a carousel protocol.

[0009] One problem with some FEC codes is that they require excessive computing power or memory to operate. Another problem is that the number of output symbols must be determined in advance of the coding process. This can lead to inefficiencies if the loss rate of packets is overestimated, and can lead to failure if the loss rate of packets is underestimated.

[0010] For traditional FEC codes, the number of possible output symbols that can be generated is of the same order of magnitude as the number of input symbols the content is partitioned into. Typically, but not exclusively, most or all of these output symbols are generated in a preprocessing step before the sending step. These output symbols have the property that all the input symbols can be regenerated from any subset of the output symbols equal in length to the original content or slightly longer in length than the original content.

[0011] “Chain Reaction Coding” as described in U.S. Patent No. 6,307,487 entitled “Information Additive Code Generator and Decoder for Communication Systems” (hereinafter “Luby I”) and in U.S. Patent Application No. 10/032,156 entitled “Multi-Stage Code Generator and Decoder for Communication Systems” (hereinafter “Raptor”) represents a different form of forward error-correction that addresses the above issues. For chain reaction codes, the pool of possible output symbols that can be generated is orders of magnitude larger than the number of the input symbols, and a random output symbol from the pool of possibilities can be generated very quickly. For chain reaction codes, the output symbols can be generated on the fly on an as needed basis concurrent with the sending step. Chain reaction codes have the property that all input symbols of the content can be regenerated from any subset of a set of randomly generated output symbols slightly longer in length than the original content.

[0012] Other descriptions of various chain reaction coding systems can be found in documents such as U.S. Patent Application No. 09/668,452, filed September 22, 2000 and entitled "On Demand Encoding With a Window" and U.S. Patent Application No. 09/691,735, filed October 18, 2000 and entitled "Generating High Weight Output symbols Using a Basis."

[0013] Some embodiments of a chain reaction coding system consist of an encoder, and a decoder. Data may be presented to the encoder in the form of a block, or a stream, and the encoder may generate output symbols from the block or the stream on the fly. In some embodiments, for example those described in Raptor, data may be pre-encoded off-line using a static encoder, and the output symbols may be generated from the plurality of the original data symbols and the static output symbols.

[0014] In some embodiments of a chain reaction coding system, the encoding and the decoding process rely on a weight table. The weight table describes a probability distribution on the set of source symbols. That is, for any number W between 1 and the number of source symbols, the weight table indicates a unique probability $P(W)$. It is possible that $P(W)$ is zero for substantially many values of W , in which case it may be desirable to include only those weights W for which $P(W)$ is not zero.

[0015] In some embodiments of a chain reaction coding system the output symbols are generated as follows: for every output symbol a key is randomly generated. Based on the key, a weight W is computed from the weight table. Then a random subset of W source symbols is chosen. The output symbol will then be the XOR of these source symbols. These source symbols are called the neighbors or associates of the output symbol hereinafter. Various modifications and extensions of this basic scheme are possible and have been discussed in the above-mentioned patents and patent applications.

[0016] Once an output symbol has been generated, it may be sent to the intended recipients along with its key, or an indication of how the key may be regenerated. In some embodiments, many output symbols may make up one transmission packet, as for

example described in the U.S. Patent Application No. 09/792,364, filed February 22, 2001 and entitled "Scheduling of multiple files for serving on a server."

[0017] In certain applications it may be preferable to transmit the source symbols first, and then to continue transmission by sending output symbols. Such a coding system is referred to herein as a systematic coding system. On the receiving side, the receiver may try to receive as many original input symbols as possible, replace the input symbols not received by one or more output symbols and use them to recover the missing input symbols. The transmission of output symbols may be done proactively, without an explicit request of the receiver, or reactively, i.e., in response to an explicit request by the receiver. For example, for applications where no loss or a very small amount of loss is anticipated, it might be advantageous to send the original input symbols first, and to send additional output symbols only in case of erasures. This way, no decoding needs to be performed if there were no losses. As another application, consider the transmission of a live video stream to one or more recipients. Where there is expectation of some loss, it may be advantageous to protect the data using chain reaction coding. Because of the nature of a live transmission, the receiver may be able to buffer a specific part of the data only for at most a pre-determined amount of time. If the number of symbols received after this amount of time is not sufficient for complete reconstruction of data, it may be advantageous in certain applications to forward the parts of the data received so far to the video player. In certain applications, and where appropriate source coding methods are used, the video player may be able to play back the data in a degraded quality. In general, where applications may be able to utilize even partially recovered data, it may be advantageous to use a systematic coding system.

[0018] Straightforward modifications of embodiments of chain reaction coding systems as described in Luby I or Raptor to produce systematic coding systems generally leads to inefficiencies. For example, if in a chain reaction coding system the first transmitted symbols comprise the original symbols, then it may be necessary to receive a number of pure output symbols which is of the same order of magnitude as the original symbols in order to be able to recover the original data. In other words, reception of the original

symbols may only minimally help the decoding process, so that the decoding process has to rely entirely on the other received symbols. This leads to an unnecessarily high reception overhead.

[0019] What is therefore needed is a systematic version of a chain reaction coding system, which has efficient encoding and decoding algorithms, and has a similar reception overhead as a chain reaction coding system.

Summary of the Invention

[0020] The present invention provides systems and methods for encoding and decoding data using systematic chain reaction encoding and decoding processes. These present can be used in numerous applications, one being a data communication system in which data is communicated faster, more reliably, and with less computational expense.

[0021] In one embodiment of the present invention, a method of encoding data into a chain reaction code is presented. Initially a set of input symbols is generated from the data. Subsequently, one or more non-systematic output symbols are generated from the set of input symbols, each of the one or more non-systematic output symbols being selected from an alphabet of non-systematic output symbols, and each non-systematic output symbols generated as a function of one or more of the input symbols. As a result of this encoding process, any subset of the set of input symbols is recoverable from (i) a predetermined number of non-systematic output symbols, or (ii) a combination of (a) input symbols which are not included in the subset of input symbols that are to be recovered, and (b) one or more of the non-systematic output symbols.

[0022] Additional embodiments and features of the invention will be better understood in view of the following drawings and detailed description.

Brief Description of the Drawings

[0023] Figs.1A and 1B illustrate exemplary embodiments of a non-systematic chain reaction encoder and decoder, respectively.

Fig. 2 illustrates the generation of an output symbol from the original input symbols used in the chain reaction decoding process.

Fig. 3 illustrates an exemplary decoding graph used in the chain reaction decoding process.

Fig. 4 illustrates a decoding matrix for the decoding graph shown in Fig. 3.

Fig. 5 illustrates an exemplary procedure for obtaining a modified decoding graph used in the chain reaction decoding process.

Fig. 6 illustrates a modified decoding equation used in the chain reaction decoding process.

Fig. 7A illustrates an exemplary method for encoding data using systematic chain reaction codes in accordance with the present invention.

Fig. 7B illustrates an exemplary method for decoding systematic chain reaction codes in accordance with the present invention.

Fig. 7C illustrates a block diagram of a communications system employing systematic coding and decoding in accordance with one embodiment of the present invention.

Fig. 8A illustrates the operation of the systematic encoder in accordance with one embodiment of the present invention.

Fig. 8B illustrates the operation of the systematic decoder in accordance with one embodiment of the present invention.

Fig. 9A illustrates one embodiment of the systematic encoder in accordance with the present invention.

Fig. 9B illustrates one embodiment of the systematic decoder in accordance with the present invention.

Fig. 10 illustrates one method for generating the systematic keys in accordance with the present invention.

Fig. 11 illustrates a second method for generating the systematic keys in accordance with the present invention.

Fig. 12 illustrates a third method for generating the systematic keys in accordance with the present invention.

Fig. 13 illustrates a forth method for generating the systematic keys in accordance with the present invention.

Fig. 14 illustrates a method for decoding a chain reaction code having systematic and non-systematic symbols in accordance with the present invention.

Figs. 15-17 illustrate the encoding and decoding processes in an exemplary embodiment of the present invention.

[0024] For clarity and convenience, features and components which are identified in earlier drawings retain their reference numerals in subsequent drawings.

Detailed Description of Exemplary Embodiments

I. Non-systematic Chain Reaction Encoder & Decoder

[0025] Figs. 1A and 1B depict exemplary embodiments of a non-systematic chain reaction encoder 130 and decoder 170, respectively, as described in Luby I and Raptor. While not referred to as such in Luby I and Raptor, these embodiments are referred to herein as “non-systematic” to differentiate their architecture and operation from the systematic encoders and decoders presented below.

[0026] Referring now to Fig.1A, the non-systematic encoder 130 accepts as input symbols $IS(0)$, $IS(1)$, ..., and keys I_0 , I_1 , ... generated by key generator 120. The number of input symbols may or may not be known in advance. In some embodiments, the non-systematic encoder 130 generates for each key I an output symbol. In Fig. 1A the output are denoted $B(I_0)$, $B(I_1)$, ... corresponding to the keys I_0 , I_1 , The number of generated output symbols is potentially limitless. Key generator 120 may have access to a random number generator from which it generates the keys. Alternatively, the keys I may be generated by some other mechanism. Encoder 130 may include static and dynamic encoders, as described for example in Raptor. It may have access to an additional key generator used to describe a static encoder.

[0027] There are various methods for obtaining the output symbols from the input symbols for which reference is made to Luby I and Raptor. One illustrative embodiment of such an encoding method is given in Fig. 2. It describes the generation of an output symbol 270 from the original input symbols. The original input symbols are denoted 210(a)-210(f). In some embodiments the first step of the coding process is static encoding, as described in Raptor. This step may produce the source symbols, denoted 220(a)-220(f), and 260(a)-260(c). In some embodiments, static encoding may be systematic, so that the values of the source symbols 220(a)-220(f) are equal to those of 210(a)-210(f). In some embodiments, there may be no static encoding, in which case the input symbols coincide with the source symbols. The source symbols can be produced off-line, or on-line, as the data symbols become available.

[0028] Once the source symbols have been created, the output symbols are generated from the source symbols. In some embodiments, the output symbol's value is the XOR of the values of some of the source symbols. For each output symbol, key generator 120 produces a key, from which the weight of the output symbol is determined from a weight table 250. Once the weight W is determined, W random or pseudorandom source symbols are chosen, and the value of the output symbol is computed as the XOR of the values of these source symbols. For example, in Fig. 2, the weight of the output symbol 270 is equal to 3 and its value is determined as the XOR of the source symbols 220(a), 220(d), and 260(b). The weight of an output symbol will also sometimes be referred to as the degree of the output symbol in this document. If a source symbol S contributes to the value of an output symbol O , then S and O are called neighbors. For example, in the situation depicted in Fig. 2 the output symbol 270 is a neighbor of each of the source symbols 220(a), 220(b), and 220(d).

[0029] Various embodiments of the chain reaction decoder 170 of Fig. 1B are described in detail in Luby I and Raptor. In some embodiments the decoding process starts as soon as enough output symbols have been collected. In some embodiments the number of collected output symbols is slightly larger than the number of original input symbols. In other embodiments, the number of collected output symbols needed to start

the decoding process can be significantly smaller than the number of original input symbols.

[0030] In some embodiments, for each received output symbol key regenerator 160 calculates the corresponding key for the symbol, and from the key determines the neighboring source symbols.

[0031] One possible description of an embodiment of a decoding process for a chain reaction decoding can be described in terms of the corresponding Decoding Graph, as exemplified in Fig. 3. This graph consists of two sets of nodes, the source nodes, and the output nodes, corresponding to the source symbols and to the received output symbols, respectively. The source nodes correspond to the source symbols, and similarly the output nodes correspond to output symbols. An output node is connected to a source node if the source symbol corresponding to the source node is a neighbor of the output symbol corresponding to the output node. In this case said output node and said source node are called neighbors.

[0032] In some embodiments the decoding starts by identifying an output node O_1 of degree one. Then the unique neighbor of O_1 is declared recovered and is removed from the Decoding Graph, and the process is continued by identifying another output node O_2 of degree one. For example, in the situation depicted in Fig. 3, O_1 could be the output node denoted 330(a). Removal of its unique neighbor, 320(b), from the Decoding Graph, leads to another output node of degree one, namely 330(c). The process continues until all the source nodes are recovered, or until there are no output node of degree one left.

[0033] For example, in the situation of Fig. 3, the following sequence of output nodes may be chosen to recover the corresponding source nodes:

Output node	Recovered source node
330(a)	320(b)

330(c)	320(e)
330(h)	320(h)
330(d)	320(i)
330(b)	320(a)
330(j)	320(f)
330(e)	320(c)
330(f)	320(g)
330(g)	320(d)

In this case decoding is successful.

[0034] In some embodiments, the graph interpretation may be used to set up a schedule for the actual computations required for the decoding, as illustrated in Luby I or Raptor. Moreover, the idealized decoder described above may be changed in a variety of ways to reduce the resources required, and to speed up the decoding process, as described in the above mentioned patents and patent applications.

[0035] In some embodiments, the decoder may output the sequence of output nodes that were used to recover the corresponding input nodes. For example, in the case outlined above, the decoder may output the indices corresponding to the output nodes 330(a), 330(c), 330(h), 330(d), 330(i), 330(b), 330(j), 330(e), 330(f), and 330(g).

[0036] It is sometimes advantageous to consider a matrix representation of the Decoding Graph, and an interpretation of the decoding algorithm in terms of this matrix, called the Decoding Matrix hereinafter. In some embodiments of the present invention the Decoding Matrix corresponding to the Decoding Graph has as many rows as there are output nodes, and as many columns as there are source nodes, and has entries 0 or 1. There is a 1 at position (k,j) of the Decoding Matrix if the j-th source node is a neighbor of the k-th output node.

[0037] Fig. 4 is an illustration of the Decoding Matrix for the Decoding Graph of Fig. 3. As is known to those skilled in the art, the decoding problem can be phrased in terms of solving a system of equations given by the Decoding Matrix. If M denotes the Decoding Matrix corresponding to the Decoding, and if the vector of values of the output symbols is denoted by b , and if there are K source nodes, then the unknown source symbol values x_1, x_2, \dots, x_K satisfy the matrix equation:

$$M \cdot x = b,$$

where x is the column vector (x_1, x_2, \dots, x_K) . The chain reaction decoding is successful if there is a permutation of rows and columns of M such that the resulting matrix is a lower triangular matrix, i.e., such that the values in the matrix above the main diagonal are zero. For example, by performing the permutation $(3 \rightarrow 2, 8 \rightarrow 3, 2 \rightarrow 5, 10 \rightarrow 6, 5 \rightarrow 7, 6 \rightarrow 8, 7 \rightarrow 9)$ on the rows, and the permutation $(2 \rightarrow 1, 5 \rightarrow 2, 8 \rightarrow 3, 9 \rightarrow 4, 1 \rightarrow 5, 3 \rightarrow 7, 7 \rightarrow 8, 4 \rightarrow 9)$ on the columns of M a lower triangular matrix is produced. Stated in terms of matrices, this means that the chain reaction decoding algorithm produces permutation matrices P and Q such that $P \cdot M \cdot Q$ is a lower triangular matrix. There are various methods for solving a system of linear equations, as is known to those of skill in the art. For example, it is possible to use the Gaussian elimination algorithm.

[0038] The matrix view of the decoding is for illustrative purposes only and not restrictive. In particular, the actual operations of the decoder may differ substantially from the preceding discussions, as described in Luby I, Raptor, and the above mentioned patent applications.

[0039] In some embodiments, if a multi-stage chain reaction coding system is used, as described in Raptor, the Decoding Graph may be augmented by a secondary graph which describes the relationship among the source symbols given by the particular static encoding used. For example, where a low-density parity-check code is used for the static encoding process, then a number of output nodes equal to the number of check symbols in this code may be added to the Decoding Graph, their value set to 0, and the Decoding Graph may be augmented by the graph of the low-density parity-check code between the source nodes and the check nodes, and the Decoding Graph may be replaced by the new

graph. The choice of low-density parity-check codes is not essential to this application. In general, for any type of static encoding, the corresponding parity-check matrix defines a bipartite graph by which the Decoding Graph may be augmented. This new graph will be referred to as the Modified Decoding Graph hereinafter.

[0040] Fig. 5 is an illustrative embodiment of a procedure for obtaining the Modified Decoding Graph. The source nodes are denoted 510(a)-510(f), the output nodes are denoted 520(a)-520(g), and the check nodes are denoted 530(a)-530(d). The source nodes correspond to the source symbols. The graph between the output nodes and the source nodes is the Decoding Graph, given by the neighborhood structure of the output nodes. The graph between the check nodes and the source nodes describes the relations among the source nodes. For example, node 530(a) shows that the XOR of the values of the source symbols corresponding to source nodes 510(a), 510(b), 510(e), and 510(f) is zero.

[0041] To the Modified Decoding Graph corresponds a Modified Decoding Matrix consisting of zeros and ones, which has as many columns as there are source nodes, and as many rows as the aggregate value of output nodes and check nodes. Correspondingly, the Modified Decoding Matrix consists of two sets of rows, one corresponding to the output nodes, and one corresponding to the check nodes. Where there are L output nodes, C check nodes, and K source nodes, the Modified Decoding Matrix may be decomposed into a submatrix M_o consisting of L rows and K columns, and a matrix M_c consisting of C rows and K columns. If x_1, \dots, x_K denote the unknown values of the source symbols, and b_1, \dots, b_L denote the known values of the received output symbols, the task of the decoder may be to solve the system of equations given by $M_o \cdot x = b$, and $M_c \cdot x = 0$. The combined system of equations would be the one given in Fig. 6.

[0042] In some embodiments of a chain reaction decoder a different decoder, called an Inactivation Decoder, may be used. This Decoder is described in greater detail in the commonly assigned co-pending US patent application no. 10/459,370, entitled "Systems and Process for Decoding a Chain Reaction Code through Inactivation," herein incorporated by reference, and referred to as the "Inactivation Decoder."

II. Systematic Chain Reaction Encoder & Decoder and Methods of Operation

[0043] Fig. 7A illustrates an exemplary method for encoding data using systematic chain reaction codes in accordance with the present invention. As used herein, the term “output symbol(s)” refers to a chain reaction code, examples of which are described in Luby I and Raptor. Systematic and non-systematic output symbols are, accordingly, specific types of chain reaction codes, a systematic output symbol comprising a transmitted input symbol, and a non-systematic output symbol comprising a output symbols which is a function of one or more input symbols.

[0044] The method of Fig. 7A may be used for a variety of applications, such as encoding data for transmission across a real-time channel, such as a path through the Internet or a broadcast link from a television transmitter to a television recipient or a telephone connection from one point to another, or the communications channel can be a storage channel, such as one or multiple CD-ROMs, disk drives, Web sites, or the like. The communications channel might even be a combination of a real-time channel and a storage channel, such as a channel formed when one person transmits an input file from a personal computer to an Internet Service Provider (ISP) over a telephone line, the input file is stored on a Web server and is subsequently transmitted to a recipient over the Internet.

[0045] Referring now to Fig. 7A, the encoding process begins at 702 when a set of input data is received, and a set of input symbols is generated therefrom. Exemplary embodiments of this process are described in Luby I and Raptor, although other techniques may be used in alternative embodiments under the present invention. As described in this document and those referred to or incorporated herein by reference, the input data may be of any format and type, including live data in which the entire set is not known a priori.

[0046] Next, one or more non-systematic output symbols are generated from the input symbols. In a particular embodiment of that process, intermediate input symbols are

initially generated from the input symbols (704). Subsequently, one or more non-systematic output symbols are generated from the intermediate input symbols (706). In alternative embodiments under the invention, the process of 706 may be omitted and the non-systematic output symbols are generated from the input symbols. Each of these processes are illustrated in greater detail below.

[0047] As will be further described below, the input symbols are provided by an input symbol generator for the input data. As explained above, the input data may be data obtained in real-time from a secondary device, such as a video capture module, or it can be static, for example when the input data resides in a file or a buffer created by a secondary application. In other applications of the present invention the input data may be acquired by a combination of a real-time and a static method, for example by receiving the data from a secondary device or application, such as a network card, and storing it on a storage device for further processing by the input symbol generator.

[0048] Fig. 7B illustrates an exemplary method for decoding systematic chain reaction codes in accordance with the present invention. Initially at 712, a first subset of the input symbols is acquired. The application will ordinarily determine how this process is accomplished. For instance, when used in a communication system, this process is performed by receiving input symbols of a chain reaction codes transmitted across a communication channel. As explained above, in particular embodiments of the present invention a communications channel can be a real-time channel, or it can be a storage channel, a combination thereof. In a particular embodiment further illustrated below, acquisition of the input symbols is accomplished by transmitting the input symbols to a receiver, the transmitted input symbols comprising systematic output symbols. Because of expected channel losses, some of the transmitted input symbols (i.e., the systematic output symbols) may be lost. Accordingly, only a subset of the original set of input symbols may be acquired by the receiver.

[0049] Next at 714, one or more non-systematic output symbols are acquired. Typically, the acquisition of the non-systematic output symbols will follow the same

modality as the input symbols, although other means may be used in alternative embodiments.

[0050] The method continues at 716, where one or more of the input symbols which were not acquired, are recovered. In a specific embodiment of this process, the missing input symbols may be recovered either from the non-systematic output symbols, or from a combination of non-systematic output symbols and the acquired input symbols.

[0051] The recovery process at 716 may be used to recover one, several, or all of the missing input symbols. Once the desired number of missing input symbols is recovered, they may be added to the acquired input symbols to re-form the original set of input symbols, and accordingly, a copy of the original data.

[0052] Fig. 7C is a block diagram of an exemplary communications system 700 that uses systematic coding and decoding in accordance with one embodiment of the present invention. In the communication system 700, an input file 721, or an input stream 725, is provided to an input symbol generator 726. Input symbol generator 726 generates a sequence of one or more input symbols (IS(0), IS(1), IS(2), ...) from the input file or stream, with each input symbol having a value and a position (denoted in Fig. 7 as a parenthesized integer). As explained above, the possible values for input symbols, i.e., its alphabet, is typically an alphabet of 2^M symbols, so that each input symbol codes for M bits of the input file. The value of M is generally determined by the use of communication system 700, but a general purpose system might include a symbol size input for input symbol generator 726 so that M can be varied from use to use. The output of input symbol generator 726 is provided to a systematic encoder 728.

[0053] The non-systematic key generator 727 generates keys I_0, I_1, I_2, \dots corresponding to the input symbols provided to the encoder 728, the non-systematic keys being used to compute the values of the non-systematic output symbols $B(I_0), B(I_1), B(I_2), \dots$ output from the encoder 728. Each non-systematic key I_0, I_1, I_2, \dots is generated so that a large fraction of the keys for the same input file are unique. In one embodiment, the non-

systematic key generator 727 comprises the key regenerator 120 illustrated in Fig. 1A above and described in Luby I and Raptor, although in other embodiments another type of device operable to generate non-systematic keys may be used.

[0054] Systematic key generator 730 generates systematic keys C_0, C_1, C_2, \dots corresponding to the input symbols provided to the encoder 728, these keys being used to recover one or more of the input symbols not received, as will be further described below. It may use random numbers generated by random number generator 735 to generate the keys. The generation of the systematic keys will be subsequently described in greater detail. The outputs of non-systematic key generator 727 and the systematic key generator 730 are provided to encoder 728.

[0055] From each non-systematic key I provided by the non-systematic key generator 727, encoder 728 generates a non-systematic output symbol, with a value $B(I)$, from the input symbols provided by the input symbol generator. The non-systematic output symbol generated may be that as described in Luby I (single stage encoding/decoding) or the output symbol described in Raptor (multiple stage encoding/decoding). The operation of an exemplary systematic encoder 728 will be described in more detail below. The value of each output symbol is generated based on its key, and on some function of one or more of the input symbols.

[0056] In some embodiments, the number K of input symbols is used by the systematic encoder 728 to select the associates. If K is not known in advance, such as where the input is a streaming file, K can be just an estimate. The value K might also be used by systematic encoder 728 to allocate storage for input symbols and any intermediate symbols generated by systematic encoder 728.

[0057] Systematic encoder 728 forwards the input symbols $IS(0), IS(1), \dots$ together with the systematic keys C_0, C_1, \dots, C_{K-1} , or an indication on how to regenerate the systematic keys to transmit module 740. When transmitted, the symbols $IS(0), IS(1), \dots$ are herein referred to as "systematic output symbols". Systematic encoder 728 may create

a copy of the input symbols for the generation of further output symbols before forwarding the input symbols to the transmit module.

[0058] Systematic encoder 728 also provides the non-systematic output symbols $B(I_0)$, $B(I_1)$, $B(I_2)$, ... to transmit module 740. Transmit module 740 is also provided the non-systematic keys (I_0, I_1, I_2, \dots) for each such output symbol from the non-systematic key generator 727. Transmit module 740 transmits the systematic and non-systematic output symbols, and depending on the keying method used, transmit module 740 might also transmit some data about the keys of the transmitted output symbols, over a channel 745 to a receive module 750. Channel 745 is assumed to be an erasure channel, but that is not a requirement for proper operation of communication system 700. Modules 740, 745 and 750 can be any suitable hardware components, software components, physical media, or any combination thereof, so long as transmit module 740 is adapted to transmit output symbols and any needed data about their keys to channel 745 and receive module 750 is adapted to receive symbols and potentially some data about their keys from channel 745. The value of K , if used to determine the associates, can be sent over channel 745, or it may be set ahead of time by agreement of encoder 728 and decoder 755.

[0059] As explained above, channel 745 can be a real-time channel, such as a path through the Internet or a broadcast link from a television transmitter to a television recipient or a telephone connection from one point to another, or channel 745 can be a storage channel, such as a CD-ROM, disk drive, Web site, or the like. Channel 745 might even be a combination of a real-time channel and a storage channel, such as a channel formed when one person transmits an input file from a personal computer to an Internet Service Provider (ISP) over a telephone line, the input file is stored on a Web server and is subsequently transmitted to a recipient over the Internet.

[0060] Receive module 750 receives the non-systematic and/or systematic output symbols from the channel 745 which it supplies to a decoder 755. Data corresponding to the keys of the received output symbols are provided to the non-systematic key regenerator 760, and the systematic key regenerator 780. In the illustrated embodiment of

Fig.7, a set of systematic output symbols denoted by $IS(x)$, $IS(y)$, ..., $IS(z)$ is received along with a set of non-systematic output symbols $B(I_a)$, $B(I_b)$, $B(I_c)$, ... In alternative embodiments, the receive module 750 may receive systematic output symbols exclusively, or a combination of systematic and non-systematic output symbols.

[0061] The non-systematic key regenerator 760 regenerates the non-systematic keys for the received non-systematic output symbols and provides these keys to the systematic decoder 755. In one embodiment, the non-systematic key regenerator 760 comprises the key regenerator 160 illustrated in Fig. 1B above and described in Luby I and Raptor, although in other embodiments another type of device operable to regenerate non-systematic keys may be used. Systematic key regenerator 180 regenerates the systematic keys C_0 , C_1 , ... and provides them to the systematic decoder 755. The systematic key regenerator 780 may have access to some shared information with the systematic key generator 730 which facilitates the regeneration of the systematic keys. Alternatively, systematic key regenerator 780 may regenerate the keys based on additional information transmitted through channel 745. In some embodiments, systematic key regenerator 780 may have access to the same random number generator 735 which may be used to generate the systematic keys. This can be in the form of access to the same physical device if the random numbers are generated on such device, or in the form of access to the same algorithm for the generation of random numbers to achieve identical behavior.

[0062] Decoder 755 uses the non-systematic keys provided by non-systematic key regenerator 760 and systematic key generator 780 together with the corresponding output symbols, to recover the input symbols (again $IS(0)$, $IS(1)$, $IS(2)$, ...). The recovered input symbols are forwarded to the input file reassembler 765. Systematic decoder 755 may forward the received systematic output symbols $IS(x)$, $IS(y)$, ..., $IS(z)$ directly to the input file reassembler 765, before recovering the remaining input symbols. In particular, if all input symbols are received, the decoder may choose to just forward the received data to input file reassembler without further computation. Input file reassembler 765 generates a copy 770 of input file 721 or input stream 725.

[0063] In the following the operations of the systematic encoder 728 and decoder 755 will be described in greater detail. In some embodiments of the present invention these units may use chain reaction encoding and decoding, as described above.

[0064] Fig. 8A illustrates the operation of the systematic encoder 728 in a specific embodiment of the invention. Initially, the systematic encoder 728 receives the input symbols $IS(0)$, $IS(1)$, ..., $IS(K-1)$ from input symbol generator 726 in Fig. 7. The input symbols may be known in their entirety at the start of the encoding, or they may only be partially known.

[0065] In this embodiment, the systematic encoder 728 has access to the non-systematic key generator 727, which generates as many non-systematic keys I_0, I_1, \dots as the number of non-systematic output symbols generated. In addition, the systematic key generator 730 generates as many systematic keys C_0, C_1, \dots, C_{K-1} as there are input symbols. Systematic Encoder 728 passes the original input symbols to the transmit module 750, these symbols being transmitted as the systematic output symbols. The systematic encoder 728 also operates to generate non-systematic output symbols $B(I_0)$, $B(I_1)$, ... for each of the keys I_0, I_1, \dots generated by non-systematic key generator 727. The operation of the systematic key generator 730 is further described below.

[0066] Systematic key generator 730 and systematic key regenerator 780 (Fig. 7) may have access to some shared information so systematic key regenerator 780 can succeed in generating the same keys as the systematic key generator 730. In some embodiments the shared information may be transmitted to the systematic key regenerator 780. In other embodiments the systematic keys may be a deterministic function of other parameters of the code, e.g., the number of input symbols and the weight table.

[0067] In some embodiments, the systematic keys may have been pre-computed for some or all relevant values of the number of input symbols. In some embodiments, the systematic keys may be re-used for different sets of input symbols. In other embodiments, the systematic keys may be re-computed for every input block, using some

shared information between the systematic key generator 730 and the systematic key regenerator 780.

[0068] Fig. 8B illustrates the operation of the systematic decoder 755 in a specific embodiment of the invention. Systematic decoder 755 receives systematic and non-systematic output symbols from receive module 750 denoted as $IS(x)$, $IS(y)$, ..., $IS(z)$, and $B(I_a)$, $B(I_b)$, ..., respectively. In a particular embodiment, systematic decoder 755 has access to the systematic key regenerator 780, and to non-systematic key regenerator 760. The output of the systematic chain reaction decoder is the set of initial input symbols $IS(0)$, $IS(1)$, ..., $IS(K-1)$.

[0069] Fig. 9A illustrates the systematic encoder 728 in more detail. The systematic encoder 728 includes a chain reaction decoder 910, and a chain reaction encoder 920. Additionally, it may have access to a memory device (not shown) to store intermediate symbols $S(0)$, $S(1)$, ..., $S(K-1)$.

[0070] Upon receiving the input symbols $IS(0)$, $IS(1)$, ..., $IS(K-1)$, and the systematic keys C_0 , C_1 , ..., C_{K-1} , chain reaction decoder 910 computes a set of intermediate input symbols $S(0)$, $S(1)$, ..., $S(K-1)$ using, for example, the decoding methods for chain reaction codes described in the patents and patent applications incorporated herein. In some embodiments of the present invention the intermediate input symbols may be stored in memory, or on disk. In other embodiments, the intermediate input symbols may be forwarded to chain reaction encoder 920 as they become available.

[0071] Chain reaction encoder 920 uses the intermediate input symbols generated by chain reaction decoder 910 together with non-systematic keys I_0 , I_1 , I_2 , ... generated by non-systematic key regenerator 727, to generate non-systematic output symbols $B(I_0)$, $B(I_1)$, In some embodiments, this encoding process may be accomplished using the input symbol encoding process described in either Luby I or Raptor, with the modification that the intermediate input symbols of the present invention are used as the input symbols of Luby I. In a particular embodiment the non-systematic output symbols

are supplied to the transmit module 140 after the input symbols $IS(0)$, $IS(1)$, ..., $IS(K-1)$. This is however not essential for the functioning of this invention. Further, the order of transmission from the transmit module 740 may vary as well.

[0072] Fig. 9B is an illustrative embodiment of the systematic decoder 755, which includes a chain reaction decoder 930, and a chain reaction encoder 940. The input to the systematic decoder includes the received output symbols some of which comprise the received systematic output symbols $IS(x)$, $IS(y)$, $IS(z)$, ..., and some of which may comprise received non-systematic output symbols $B(I_a)$, $B(I_b)$, In some embodiments, the decoder may copy the received systematic symbols to a memory device, and directly forward them to input file reassembler 765.

[0073] Chain reaction decoder 930 uses the symbols $IS(x)$, $IS(y)$, ..., $IS(z)$, $B(I_a)$, $B(I_b)$, ..., the systematic keys C_x , C_y , ..., C_z , generated by the systematic key regenerator 780, and the non-systematic keys I_a , I_b , ... generated by non-systematic key regenerator 760 to produce intermediate input symbols $S(0)$, $S(1)$, ..., $S(K-1)$. The systematic keys C_x , C_y , ..., C_z , correspond to the received input symbols $IS(x)$, $IS(y)$, ..., $IS(z)$. In some embodiments, the recovered intermediate symbols may be stored to a secondary storage before being passed to the chain reaction encoder 440. In other embodiments, these intermediate symbols may be passed directly to the chain reaction encoder 940.

[0074] Chain reaction encoder 940 uses the intermediate input symbols and the systematic keys C_u , C_v , ..., C_w corresponding to erased systematic output symbols $IS(u)$, $IS(v)$, ..., $IS(w)$ to generate and output the missing original input symbols $IS(u)$, $IS(v)$, ..., $IS(w)$. As an exemplary embodiment, for each of the initial keys C_u , C_v , ..., C_w , the decoder identifies a weight W and W symbols among the intermediate input symbols $S(0)$, ..., $S(K-1)$, and XOR's the values of output symbols to obtain the erased input symbols $IS(u)$, $IS(v)$, ..., $IS(w)$ corresponding to the systematic keys C_u , C_v , ..., C_w . The amount of computational resources used by chain reaction encoder 940, in one embodiment, will be proportional to the number of systematic output symbols that are erased. For example, if all the systematic output symbols are received, then the decoder

may not perform any computations, and forward the received symbols to input file reassembler 765.

[0075] In particular embodiments, the chain reaction encoder 940 and chain reaction decoder 910 will have access to the same weight table, and use the same static encoding/decoding, if static encoding is used. Similarly, chain reaction encoder 920 and chain reaction decoder 930 may have access to the same weight table, and use the same static encoding/decoding.

Methods for Calculating the Systematic Keys

[0076] In a specific embodiment of the present invention, the systematic keys are calculated by systematic key generator 730 before symbol transmission, and re-computed by the systematic key regenerator 780 after symbol reception. The systematic keys are used by the chain reaction decoder 910 and encoder 930 to obtain the intermediate input symbols $S(0), S(1), \dots S(K-1)$.

[0077] In particular embodiments of the present invention the systematic keys are calculated in such a way that unique and efficient chain reaction decoding of K symbols is possible using exactly K output symbols generated with these keys. Here decoding can be any of the decoding methods described in Luby I, Raptor, or Inactivation Decoding, or more generally decoding methods based on the Gaussian elimination algorithm as for example described in Inactivation Decoding.

[0078] Fig. 10 is an exemplary embodiment of the systematic key generation process. One input to the systematic key generator may be the number K of input symbols $IS(0), IS(1), \dots, IS(K-1)$. Systematic key generation starts by setting a variable j equal to 0. During the algorithm a matrix M with K columns, which, initially, has zero rows, is updated by adding rows as the algorithm progresses. For every different value of j the algorithm generates a different key $D(j)$ at 1020. This key may be generated by the methods described in Luby I or Raptor, and may use the random number generator 135 shown in Fig. 1. Next at 1030, the key $D(j)$ is used to compute the entries of the j -th row

of the matrix M . One possible embodiment of such a computation would be to use key $D(j)$ in the chain reaction coding process. In this case, using the weight table, the key $D(j)$ identifies a weight W and W values among the values $0, 1, \dots, K-1$. It then may set a 1 at position m of the j^{th} row of M if m is one of the random or pseudorandom values generated, and set the other values of the j^{th} row to zero.

[0079] At 1040, a determination is made as to whether the presently configured matrix M has K rows that are linearly independent over the binary field $GF(2)$, the binary field $GF(2)$ referring to the set consisting of 0 and 1 in which multiplication and addition are performed modulo the integer 2. This process in 1040 can be performed in a variety of ways. For example, Gaussian elimination over the binary field $GF(2)$ could be used to check this. However, there are many other ways as known to those skilled in the art. For example, if the teachings of Inactivation Decoding are applied to the matrix M , then M contains K linearly independent rows only if the Inactivation Decoder applied to M is successful.

[0080] If the test in 1040 is positive, and rows $r(0), r(1), \dots, r(K-1)$ of M are discovered to be linearly independent, then the systematic keys C_0, C_1, \dots, C_{K-1} are set to the keys $D(r(0)), \dots, D(r(K-1))$, and the keys are output. If the test in 1040 is negative, then the counter j is incremented in 1060, and the computation is repeated from 1020 on.

[0001] Other equivalent or substantially similar methods of generating the systematic keys can be envisioned by those skilled in the art. For example, instead of generating the keys $D(j)$ one at a time during the course of the algorithm, a set of L such keys could be generated beforehand, and key $D(j)$ could be taken from this pool of keys at step j of the algorithm. Herein, L could be a function of the number of input symbols.

[0081] A second method for generating the systematic keys is exemplified in Fig. 11. In this method, the input to this algorithm consists of the number K of input symbols, and a number L which is typically larger than or equal to K . In some embodiments, L may be the number of output symbols to be collected to guarantee, with high probability, that the decoding is successful, as described in Luby I or Raptor.

[0082] At 1110, L keys $D(0), \dots, D(L-1)$ are generated. This process may be accomplished through the use of a random number generator 735. In other embodiments, these keys may be generated from a fixed list of re-usable keys. This process may also provide an indication of how the keys were generated. For example, if a random number generator is used, the seed for the generator may be recorded for future use by the systematic key regenerator.

[0083] Using the keys $D(0), D(1), \dots, D(L-1)$ a Modified Decoding Graph is set up in 1120 as described above and exemplified in Fig. 5. This process may employ the knowledge of the specific weight table for the code, as well as the knowledge of any static encoding used, as described in Raptor.

[0084] At 1130, the Modified Decoding Graph is decoded using any of the methods presented earlier. As a by-product of the decoding, the indices $r(0), r(1), \dots, r(K-1)$ of those output nodes that trigger the recovery of an input node are recorded. At 1140, the systematic keys are outputted as $C_0=D(r(0)), \dots, C_K=D(r(K-1))$.

[0085] Fig. 12 illustrates a third method for computing the systematic keys. Similar to the method of Fig. 11 the keys $D(0), \dots, D(L-1)$ are generated in 1210, and the Decoding Graph is set up using these keys, and possibly the weight table. Next a set S is initialized as the empty set in 1230. The set S will contain the indices of those output symbols which are used in the chain reaction decoding process to recover the value of an input node. In 1240 the chain reaction decoding process is applied to the Decoding Graph by identifying an output node of degree one. The index of this output node is added to the set S in accordance with the above-mentioned role of this set. A test is performed at 1250 as to whether the set S already has the right number of elements. If not, the algorithm loops back to 1240 where another input node of degree one is chosen to continue the decoding process. If the size of S is K , then the elements of S are sorted starting with the smallest element to yield the sorted elements S_0, \dots, S_{K-1} and the systematic keys are calculated as $C_0=D(S_0), \dots, C_{K-1}=D(S_{K-1})$ in 1260.

[0086] Fig. 13 illustrates a fourth method for computing systematic keys in accordance with the present invention. In this method it is assumed that a decoding algorithm is available which on input K and a set of keys can decide whether the original K symbols are decodable from the given set of keys. Examples of such algorithms are provided by the decoders described in Luby I, Raptor, or Inactivation Decoding.

[0087] At 1310 L keys $D(0), \dots, D(L-1)$ are generated. Similar to the above description, this process may be accomplished through the use of a random number generator 735, or the keys may be generated from a fixed set of re-usable keys. At 1315, the decoder is used to decide whether or not it is possible to decode the K symbols from the set of keys $D(0), \dots, D(L-1)$. If decoding is not successful, then the given set of keys does not contain as a subset the systematic keys, and the algorithm aborts at 1325. Otherwise, three sets are initialized at 1330. These sets are called Systematic, Non_Systematic, and Unvisited, respectively. At the end of the algorithm, the set Systematic will contain the set of systematic keys. Originally, at 1330 the sets Systematic and Non_Systematic are initialized to empty sets, while the set Unvisited contains all the original keys $D(0), \dots, D(L-1)$. At processes 1335 through 1360 a key is removed from the set Unvisited and a decoding attempt is made on the keys contained in the sets Systematic and Unvisited. If the attempt is successful, then the chosen key C does not belong to the set of systematic keys. On the contrary, if decoding is not successful, then the key does belong to the set of systematic keys. The procedure consisting of removal of an unvisited key and decoding (1335), a test as to whether decoding was successful (1340), and the following addition of the chosen key to the set Systematic or Non_Systematic based on the outcome of the decoder (1345 and 1350) are repeated as long as the set Systematic has fewer than the number K of original input symbols.

[0088] Fig. 14 illustrates a method for decoding a chain reaction code having systematic and non-systematic symbols in accordance with the present invention. At 1410, non-systematic keys I_a, I_b, \dots corresponding to the received non-systematic output symbols $B(I_a), B(I_b), \dots$ are used to generate a matrix B which has as many rows as there are received non-systematic output symbols and as many columns as there are input

symbols. For each key the same mechanism as for encoding chain reaction codes is used to generate a weight W and a set J_1, J_2, \dots, J_W of indices of input symbols from which the output symbol corresponding to the key is generated. Then, in the corresponding row of the matrix B the positions corresponding to J_1, J_2, \dots, J_W are set to 1, while the other positions in that row are set to 0. The procedure is repeated until all keys corresponding to non-systematic received symbols are exhausted.

[0089] Next at 1420, a similar procedure is applied to construct a square matrix C with as many rows and columns as the number of input symbols from the systematic keys C_0, C_1, \dots, C_{K-1} . This process also computes the inverse of the matrix C , called A . Computing the inverse of A can be performed in a variety of ways, as is known to those of skill in the art. For example, a Gaussian elimination algorithm can be used to calculate A . In other embodiments a version of chain reaction decoding can be utilized to perform this step. This is further illustrated in an example later in this disclosure.

[0090] At 1430, the product of the matrices B and A is calculated over the binary field $GF(2)$ to obtain a matrix H . Next at 1440, two sets of indices E and R are determined: E is the set of indices of the non-received systematic symbols, while R is the set of indices of the received systematic symbols. For example, assume there are 11 input symbols with indices 0, 1, 2, ..., 10. If, after the transmission, the systematic symbols corresponding to the indices 0, 3, 9, 10 are received, then $R = \{0, 3, 9, 10\}$, while $E = \{1, 2, 4, 5, 6, 7, 8\}$. The matrix H , computed in 1430 as the product of B and A is then subdivided into two submatrices H_E and H_R : H_E is the submatrix of H obtained by taking the columns of H corresponding to the indices of the systematic symbols not received, and H_R is the submatrix of H corresponding to the indices of the received systematic symbols. In the example above, H_E would be the submatrix of H formed by the columns 1, 2, 3, 4, 5, 6, 7, and 8 of H .

[0091] At 1450, the matrix H_R is multiplied with the vector formed by the received systematic symbols $IS(x), IS(y), \dots, IS(z)$. For example, in the scenario above, H_R would be multiplied with the values of the systematic symbols 0, 3, 9, 10 (in this ordering). The

actual multiplication can be performed in a variety of ways, as is known to those skilled in the art. The result of this multiplication, called the vector y in the following, may be stored for future use. At 1460, the non-systematic received output symbols are used to set up a vector b . Where there are L such symbols, the number of entries in the vector b is L . This step may only be logical. In other words, this step may not require any computations. Next, the results of the previous multiplication stored in the vector y is component-wise XOR'd with the entries of the vector b , i.e., each of the non-systematic received output symbols are XOR'd with the corresponding symbols of the vector y . The result of this operation may be stored in place of the received non-systematic symbols, or it may be stored at a different location.

[0092] Once this XOR has been determined, a system of linear equations is set up using the matrix H_E corresponding to the erased systematic symbols. The solution x of the system $H_E * x = y + b$ then corresponds to the values of the erased systematic symbols. These values are output in 1470. Again, this process can be performed in a variety of ways, for example using Gaussian elimination, or any of the variants of chain reaction decoding disclosed in Luby I, Raptor, or Inactivation Decoding.

[0093] This matrix view of the decoding is for illustrative purposes only and not restrictive. Many variations of this decoding procedure will become apparent to those of skill in the art upon review of this disclosure.

III. Exemplary Systematic Coding and Decoding

[0094] A brief example of some aspects of the operations of some embodiments of a systematic chain reaction coding system will now be given with reference to Figs. 15-17. In all the examples given the effect of the weight table is only implicitly stated in terms of the list of neighbors of a given symbol, given its key.

Computing the Systematic Keys

[0095] Fig. 15A describes a Decoding Graph used to obtain systematic keys C_0, C_1, \dots, C_8 . It is assumed that 12 keys $D(0), D(1), \dots, D(11)$ have already been generated, for example by the operation in 1110 of Fig. 11. The graph in Fig. 15A describes the Modified Decoding Graph between the input nodes denoted 1520(a), ..., 1520(i), and output nodes denoted 1530(a), ..., 1530(l) using the keys $D(0), \dots, D(11)$. Chain reaction decoding may now be applied to this graph to obtain the systematic keys as the keys of those output nodes which trigger the recovery of an input node in the course of chain reaction decoding.

[0096] In operation, node 1530(a) may be used to recover the input node 1520(b). Accordingly, the first systematic key C_0 is then equal to the first of the generated keys, namely $D(0)$. Recovery of input node 1520(b) causes output node 1530(c) to become of degree 1, and hence to trigger recovery of node 1520(e). Continuing in this way, it can be seen that the nodes colored light gray in Fig. 15A can be used to recover the input nodes. The sequence of output nodes used to recover the input nodes is equal to 1530(a), 1530(b), 1530(c), 1530(d), 1530(e), 1530(f), 1530(g), 1530(h), 1530(j). As a result, the sequence of systematic keys may be chosen as shown in Fig. 15B.

[0097] It should be noted that the recovery process for the illustrated chain reaction decoding is only conceptual. In particular, no XOR operation is performed in this particular example.

Systematic Encoding

[0098] As outlined in Fig. 9A, a systematic chain reaction encoder consists of a chain reaction decoder 910 and a chain reaction encoder 920. Accordingly, the operation of systematic chain reaction encoding is divided into two parts. These two parts are exemplified in Fig. 16A and Fig. 16B, respectively.

[0099] Fig. 16B exemplifies the operation of the chain reaction decoder 910. The input symbols are denoted by $IS(0), \dots, IS(8)$. The keys C_0, C_1, \dots, C_8 are used to set up the

graphical dependency between the input symbols and the intermediate input symbols $S(0)$, ..., $S(8)$. For example, the key C_0 shows that $IS(0)$ is equal to the value of $S(1)$, while the key C_4 shows that $IS(4)$ is equal to the XOR of the values of $S(2)$, $S(5)$, and $S(7)$. Chain reaction decoding can now be applied to obtain the values $S(0)$, $S(1)$, ..., $S(8)$. The schedule to obtain these values may have been forwarded to the chain reaction decoder 910 from the systematic key generator 730 in Fig. 7, since this schedule was set up to obtain the keys C_0 , C_1 , ..., C_8 . Unlike the operation of the systematic key generator, this step may employ XOR'ing the values of the individual symbols.

[0100] In the example of Fig. 16A the schedule may first produce the value of $S(1)$, which in turn may produce the value of $S(4)$ using the value of $IS(1)$. This triggers the recovery of the values of $S(0)$, and $S(7)$, etc.

[0101] Fig. 16B exemplifies the operation of the chain reaction encoder 920 in Fig. 9A by showing the generation of the first 11 non-systematic output symbols $O(0)$, ..., $O(10)$. (The illustrated output symbols $O(i)$ refers to previously described output symbols $B(I_i)$.) As was described before, the output of the systematic encoder consists of the systematic output symbols $IS(0)$, ..., $IS(8)$, followed by the output symbols $O(0)$, ..., $O(10)$, This particular ordering is only exemplary, and other orderings can be used in alternative embodiments under the present invention.

Systematic Decoding

[0102] Figs. 17A and 17B exemplify an embodiment of the process of systematic chain reaction decoding. It is assumed that the received systematic output symbols are $IS(1)$, $IS(6)$, and $IS(7)$, while the received non-systematic output symbols are $O(0)$, $O(3)$, $O(4)$, $O(6)$, $O(7)$, $O(8)$, $O(9)$, and $O(10)$. The task of the decoder is to compute the values of the missing systematic output symbols, i.e., the values $IS(0)$, $IS(2)$, $IS(3)$, $IS(4)$, $IS(5)$, and $IS(8)$. Fig. 17A is an example of how the chain reaction decoder 930 and the chain reaction encoder 940 in Fig. 9B may be combined into one decoder. In some applications, such a combination may lead to computational savings.

[0103] Using the keys C_1 , C_6 , and C_7 corresponding to the received systematic output symbols, and the keys corresponding to the received non-systematic output symbols, a graph is set up between the received output symbols, and the intermediate input symbols $S(0)$, ..., $S(8)$. A connecting line is drawn between an output symbol and all the intermediate input symbols whose XOR yields the value of the output symbol. The individual connections are the same as the ones shown in Fig. 16A and Fig.16B. The particular ordering of the received output symbols may not be equal to the ordering chosen to represent the Decoding Graph.

[0104] This graph is extended by another layer of nodes, corresponding to the erased systematic output symbols. This graph corresponds to the upper part of Fig. 17A, in which the input symbols $IS(0)$, $IS(2)$, $IS(3)$, $IS(4)$, $IS(5)$, and $IS(8)$ are connected via dotted lines to those intermediate input symbols of which they are an XOR of. Again, these connections may be verified against the corresponding connections in Fig. 17A.

[0105] The process of decoding in this particular example may start by applying the chain reaction decoding to the lower graph; every time one of the intermediate symbols is recovered, its value may be XOR'd to the value of the all the neighbors of this symbol among the non-received original symbols in the upper part of the figure.. Originally, the values of these symbols may be set to zero.

[0106] For example, output symbol $O(4)$ may be used to recover the value of $S(3)$. The value of $S(3)$ may then be XOR'd into the current value of $IS(5)$. After this step, the value of $IS(5)$ is equal to that of $S(3)$. Recovery of $S(3)$ reduces the degree of the output node $O(10)$ to one. This output node in turn recovers the value of the intermediate symbol $S(6)$. This value is XOR'd into the current value of $IS(5)$, so that after this step the value of $IS(5)$ is recovered. The process may continue until all the non-received systematic input symbols are recovered.

[0107] Fig. 17B illustrates the process by which the missing output symbols are recovered. The recovered symbols are framed in rectangles. The recovered systematic

output symbols are framed in gray rectangles. The labels of the edges in this figure describe the symbols used for the recovery.

[0108] For example, symbol O(4) is used to recover S(3). Symbol O(10) is used to recover S(6). S(3) and S(6) together recover S(5). Recovery of S(6) triggers the recovery of S(8) (using O(9)) and the recovery of S(0) (using the received systematic output symbol IS(7)). Recovery of S(8) triggers the recovery of IS(3). Recovery of S(0) triggers the recovery of S(4) (using IS(1)). On the other hand, using O(0), the recovery of S(8) triggers that of S(1), which together with S(4) recovery IS(2). Furthermore, recovery of S(1) leads to recovery of IS(0), since these values are identical. Using O(8), and the recovered value of S(4), the value of S(5) is obtained. This, in turn, recovers the value of IS(8), since the latter is the XOR of S(5), S(4), and S(0), and all these values are known at this stage. Using IS(6) and S(4), the value of S(7) is obtained. Using O(7), this recovers the value of S(2), which together with S(7) recovers the value of the last remaining input symbol, namely IS(4).

[0109] The foregoing description has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously many modifications and variations are possible in light of the above teaching. The described embodiments were chosen in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

Documents Herein Incorporated by Reference:

[0110] U.S. Patent No. 6,307,487 to Michael G. Luby, entitled "Information Additive Code Generator and Decoder for Communication Systems" (referred to herein as Luby I);

U.S. Patent Application No. 09/792,364, filed February 22, 2001, entitled "Scheduling of Multiple Files for Serving on a Server";

U.S. Patent Application No. 10/032,156, filed December 21, 2001, entitled “Multi-Stage Code Generator and Decoder for Communication Systems” (referred to herein as “Raptor”); and

U.S. Patent Application No. 10/459,370, filed June 10, 2003, entitled “Systems and Processes for Decoding Chain Reaction Codes through Inactivation” (referred to herein as “Inactivation Decoding”).